

AcSel: selecting variables with accuracy in correlated data sets

Nicolas Jung ^{*1,2}, Frédéric Bertrand¹, and Myriam Maumy-Bertrand¹

¹Institut de Recherche en Mathématique Avancée, UMR 7501

²Laboratoire d'Immunogénétique Moléculaire Humaine, Institut National de la Santé et de la Recherche Médicale, Unité Mixte de Recherche S1109

December 11, 2015

Abstract

With the emergence high-throughput technologies, it is possible to measure large amounts of data relatively low cost. Such situations arise in many fields from sciences to humanities, and variable selection may be of great help to answer challenges that are specific to each of them. Variable selections may allow to know, among all measured variables, which are of interest and which are not. A lot of methods have been proposed to handle this issue, with the Lasso and other penalized regression as special cases. These methods fail in some cases and linear correlation between explanatory variables is the most common, especially in big datasets. In this article, we introduce AcSel, a wrapping algorithm to enhance the accuracy of any variable selection method.

1 Introduction

The problem of variable selection has received an increasing attention over the last years [FL06] and is one of the most important challenges for the 21st century [D⁺00]. Indeed, technological innovations make it possible to measure large amounts of data relatively low cost. As a consequence, problems in which the number P of variables is greater than the number N of observations have become common. As reviewed by Fan and Li [FL06], such situations arise in many fields from sciences to humanities, and variable selection may be of great help to answer challenges that are specific to each of them. For example, in biology, thousands of messenger RNA (mRNA) gene expressions [LFGL99] may be potential predictors of some illness. Other examples are imagery (magnetic resonance image, nuclear magnetic resonance, satellite images...), financial engineering and risk management, or health studies [FL06]. Moreover, in such studies, the correlation between variables is often very strong [SDC03] and variable selection methods often fail to choose the informative variables among those which are not.

In this article, we assume that our data is generated by a multivariate linear model:

*Corresponding author, please write at : nicolas.julien.jung@gmail.com

$$\mathbf{y} = \mu \mathbf{1}_N + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (1)$$

where $\mathbf{y} = (y_1, \dots, y_N)'$ is the response variable, μ is the mean variable response, $\mathbf{1}_N$ is a vector of length N containing only ones, $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_P)$ is the design matrix of size $N \times P$, $N > 2$, with $\mathbf{x}_p = (x_{p1}, \dots, x_{pN})'$ which are the variables and $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_N)$ is a Gaussian noise vector which is the realization of some random law with a mean of 0 and an unknown variance σ^2 . Furthermore, we will assume that the vector of parameters $\boldsymbol{\beta} = (\beta_1, \dots, \beta_P)'$ is sparse. In other words, we will assume that $\beta_i = 0$ except for a quite small proportion of elements of the vector. We note \mathcal{S} as the set of indexes for which $\beta_i \neq 0$ and $q < \infty$ is the cardinal of this set \mathcal{S} . Without any loss of generality, we will assume that $\beta_p \neq 0$ if and only if $p \leq q$. Moreover, we assume that the response and the variables are centred and that $\|\mathbf{x}_p\|^2 = 1$ for $p = 1, \dots, P$ where $\|\cdot\|$ stands for the usual euclidean norm; in this context, we have $\mu = 0$.

When dealing with a problem of variable selection, there are three main goals. We enumerate them in increasing level of difficulty:

1. The prediction goal, in which you want $\hat{\mathbf{y}}$ to be as close as possible to \mathbf{y} .
2. The estimation goal, in which you want $\hat{\boldsymbol{\beta}}$ to be as close as possible to $\boldsymbol{\beta}$.
3. The estimation of the support, in which you want $\mathbb{P}(\mathcal{S} = \hat{\mathcal{S}})$ to be close to one.

Fan and Li [FL01] proposed another desirable property, the oracle property, which combines goals 2 and 3. Precisely, a method is said to have the oracle property if it discovers the correct support, and if the rate of convergence of $\hat{\boldsymbol{\beta}}$ toward $\boldsymbol{\beta}$ is optimal (*i.e.* the same as in the case in which the correct support is known). Here, our interest is mainly in the third goal, *i.e.* in identifying the correct support \mathcal{S} . This kind of issue arises in many fields, for example in biology, where it is of greatest interest to discover which specific molecules are involved in a disease [FL06].

There is a vast literature dealing with the problem of variable selection in both statistical and machine learning areas ([FL06, FL10]). The main variable selection methods can be gathered in the common framework of penalized likelihood. The estimate $\hat{\boldsymbol{\beta}}$ is then given by:

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^P} \left[-\ell_N(\boldsymbol{\beta}) + \sum_{p=1}^P \text{pen}_{\boldsymbol{\lambda}}(\beta_p) \right], \quad (2)$$

where $\ell_N(\cdot)$ is the log-likelihood function, $\text{pen}_{\boldsymbol{\lambda}}(\cdot)$ is a penalty function with k parameters and $\boldsymbol{\lambda} = (\lambda, \lambda_2, \lambda_3, \dots, \lambda_k)'$. As the goal is to obtain a sparse estimation of the vector of parameters $\boldsymbol{\beta}$, a natural choice for the penalty function is to use the so-called L_0 norm ($\|\cdot\|_0$) which corresponds to the number of non-vanishing elements of a vector:

$$\begin{aligned} \text{pen}_{\boldsymbol{\lambda}} &: \mathbb{R} \rightarrow \{0, \lambda\} \\ x &\mapsto \begin{cases} \text{pen}_{\boldsymbol{\lambda}}(x) = \lambda & \text{if } x \neq 0 \\ \text{pen}_{\boldsymbol{\lambda}}(x) = 0 & \text{else} \end{cases} \Rightarrow \sum_{p=1}^P \text{pen}_{\boldsymbol{\lambda}}(\beta_p) = \lambda \|\boldsymbol{\beta}\|_0. \end{aligned} \quad (3)$$

For example, when $\lambda = 1$, we get the Akaike Information Criterion (AIC) [Aka74] and when $\lambda = \frac{\log(N)}{2}$ we get the Bayesian Information Criterion (BIC) [Sch78]. Another slightly different

formulation leads to Mallows's C_p [Mal73] or to the Risk Inflation Criterion [FG94]. In the context of Gaussian independent and identically distributed (i.i.d.) errors in the model described in equation (1), the following holds [BA02]:

$$-\ell_N(\boldsymbol{\beta}) = \frac{N}{2} \log \left(\frac{\|\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}\|^2}{N} \right) + K_1, \quad (4)$$

where K_1 is a constant. Up to an affine transformation of the log-likelihood [FL10], we see that equation (2) is equivalent to:

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^P} \left[\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2 + \sum_{p=1}^P \text{pen}_{\boldsymbol{\lambda}}(\beta_p) \right]. \quad (5)$$

A lot of different penalties can be found in the literature. Solving this problem with $\|\cdot\|_0$ as part of the penalty is an NP-hard problem [Nat95, FL10]. It cannot be used in practice when P becomes large, even when it is employed with some search strategy like forward regression, stepwise regression [Hoc76], genetic algorithms [KBIS99]... Donoho and Elad [DE03] show that relaxing $\|\cdot\|_0$ to norm $\|\cdot\|_1$ ends, under some assumptions, to the same estimation. This result encourages the use of a wide range of penalty based on different norms. For example, the case where $\text{pen}_{\boldsymbol{\lambda}}(\beta_p) = \lambda|\beta_p|$ is the Lasso estimator [Tib96] (or equivalently Basis Pursuit Denoising [CDS01]) whereas $\text{pen}_{\boldsymbol{\lambda}}(\beta_p) = \lambda\beta_p^2$ leads to the Ridge estimator [HK70]. These two last cases can be seen as a special case of Bridge regression [FF93] in which $\text{pen}_{\boldsymbol{\lambda}}(\beta_p) = \lambda|\beta_p|^b$ with $0 < b \leq 2$. Nevertheless, the penalty term induces variable selection only if:

$$\min_{x \geq 0} \left(\frac{d\text{pen}_{\boldsymbol{\lambda}}(x)}{dx} + x \right) > 0.$$

This explains why the Lasso regression allows variable selection while the Ridge regression does not. As it is well known [Zou06a], the Lasso leads to a biased estimate. The SCAD (smoothly clipped absolute deviation) [Fan97], MCP (minimax concave penalty) [Zha10] or adaptative Lasso [Zou06a] penalties all address this problem. The popularity of such variable selection methods is linked to fast algorithms including LARS (least-angle regression) [EHJT04], coordinate descent [WL08] or PLUS [Zha10].

Nevertheless, the goal of identifying the correct support of the regression is complicated and the reason why variable selection methods fail to select the set of non-zero variables \mathcal{S} can be summed up in one word: linear correlation. Choosing the Lasso regression as a special case, Zhao and Yu [ZY06] (and simultaneously Zou [Zou06b]) found an almost necessary and sufficient condition for Lasso sign consistency (*i.e.* selecting the non-zero variables with the correct sign). This condition is known as “irrepresentable condition”:

$$\left| \mathbf{X}'_{\setminus \mathcal{S}} \mathbf{X}_{\mathcal{S}} (\mathbf{X}'_{\mathcal{S}} \mathbf{X}_{\mathcal{S}})^{-1} \text{sgn}(\boldsymbol{\beta}_{\mathcal{S}}) \right| < \mathbf{1}, \quad (6)$$

where $\mathbf{X}_{\mathcal{S}} = (x_{ij})_{i,j \in \mathcal{S}^2}$, $\boldsymbol{\beta}_{\mathcal{S}} = (\beta_p)_{p \in \mathcal{S}}$. In other words, when $\text{sgn}(\boldsymbol{\beta}_{\mathcal{S}}) = \mathbf{1}_q$, this can be seen as the regression of each variable which is not in \mathcal{S} over the variables which are in \mathcal{S} . As all variables in the matrix \mathbf{X} are centred, the absolute sum of the regression parameters should be smaller than 1 to satisfy this “irrepresentable condition”.

Facing this issue, existing variable selection methods can be split into two categories:

- Those which are “regularized” and try to give a similar coefficients to variables which are correlated (*e.g.*: elastic net [ZH05]),
- Those which are not “regularized” and pick up one variable among a set of correlated variables (*e.g.*: the Lasso [Tib96]).

The former group can then be split into methods in which groups of correlation are known, such as the group Lasso [YL06, FHT10] and those in which groups are not known as in the elastic net [ZH05]. The latter combines the \mathcal{L}_1 and the \mathcal{L}_2 norm and takes advantage of both. Broadly speaking, non-regularized methods will select some co-variables among a group of correlated variables while regularized methods will select all variables in the same group with similar coefficients (see example in Figure 2).

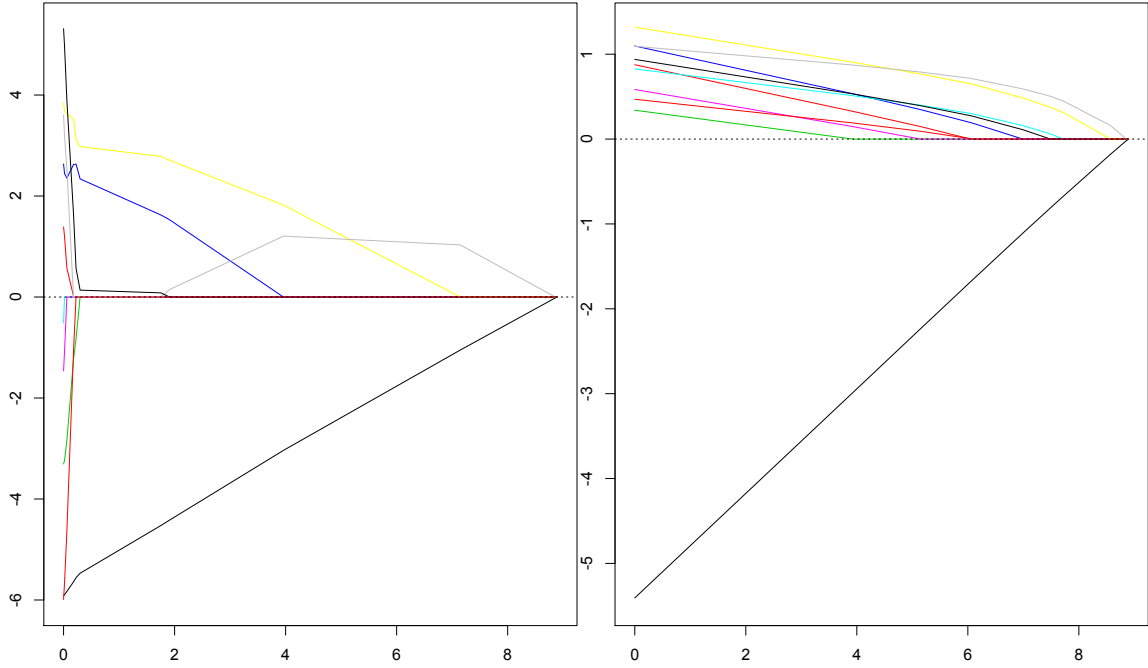


Figure 1: In this example $N = 20, P = 10, \beta = (1, 1, 0, \dots, 0)'$. The mean correlation between x_1 and the other variables is 0.20 while the mean correlation between all the other variables is 0.95. The x-axis corresponds to the value of the penalty parameter λ ; the greater the parameter, the stronger the constraint. Right: with the lasso regression, no regularization is made. Left: with the elastic net regression, the coefficients of correlated variable are similar.

However, none of these selection methods distinguishes between variables that were selected for inclusion in the model with confidence and those that were not. In this article, we propose the AcSel algorithm that can provide a confidence factor for selected variables. Our new algorithm

will be useful in different contexts, including biology where it will allow high precision selection of relevant therapeutic targets.

The rest of this article is organized as follows. In section 2 we present our new algorithm, in section 3 we drive some simulation studies. A real dataset will be analysed in section 4, while section 5 will end with some remarks and conclusion notes.

2 Methods

The AcSel algorithm has been designed in a general framework whose goal is to enhance the abilities of any variable selection method, especially those which are not regularized¹. The main goal of this algorithm is to improve the precision, *i.e.* the proportion of selected variables which really are in S .

2.1 Introduction

The main idea of our algorithm is to consider that groups of variables of the matrix \mathbf{X} which are linearly correlated are independent realizations of the same random function. According to this random function, correlated variables are then perturbed. Strictly speaking, the use of noise to make the difference between the informative and the uninformative variables is not a new idea. For example, it has been shown that adding random pseudo-variables decreases over-fitting [WBS07]. In the case where $P > N$ the pseudo-variables are generated either with independent Gaussian laws $\mathcal{N}(0, 1)$ or by using permutations on the matrix \mathbf{X} [WBS07]. Another approach consists in adding noise to the response variable and leads to similar results [LSB06]. The rational of this last method is based on the work of Cook and Stefanski [CS94] which introduces the simulation-based algorithm SIMEX [CS94]. Adding noise to the matrix \mathbf{X} has already been used in the context of microarrays [CGHE07]. Simesel [EZ12] is an algorithm that both adds noise to variables and uses random pseudo-variables. One new and interesting approach is stability selection [MB10] in which the variable selection method is applied on sub-samples, and informative variables are defined as variables which have a high probability of being selected. Bootstrapping has been applied to the Lasso on both response variable and the matrix \mathbf{X} with better results in the former case [Bac08]. The random Lasso, in which variables are weighted with random weights, has also been introduced [WNRZ11].

In this article, following the idea of using simulation to enhance the variable selection methods, we propose the AcSel algorithm. Unlike other algorithms reviewed above, our algorithm takes care of the correlation structure of the data. Furthermore, our algorithm is motivated by the fact that in the case of non-regularized variable selection methods, if a group contains variables that are highly correlated together, one of them will be chosen “at random” [ZH05].

As we assume that the variables are centred and that $\|\mathbf{x}_p\|^2 = 1$ for $p = 1, \dots, P$, we know that $\mathbf{x}_p \in \mathcal{S}^{N-2}$. Indeed, the normalization puts the variables on the unit sphere \mathcal{S}^{N-1} . The process of centring can be seen as a projection on the hyperplane \mathcal{H}^{N-1} with the unit vector as

¹Regularized should be understood as it is defined in this article, namely, a regularized method should give similar estimations for coefficients if the corresponding variables are strongly correlated.

normal vector. Moreover, the intersection between \mathcal{H}^{N-1} and \mathcal{S}^{N-1} is \mathcal{S}^{N-2} . We further define the following isomorphism:

$$\begin{aligned} \phi &: \mathcal{H}^{N-1} \rightarrow \mathbb{R}^{N-1} \\ \mathbf{h}_n &\mapsto \phi(\mathbf{h}_n) = \mathbf{f}_n \quad n = 1, \dots, N-1, \end{aligned} \quad (7)$$

where $\{\mathbf{h}_n\}_{n=1, \dots, N-1}$ is an orthogonal base of \mathcal{H}^{N-1} and $\{\mathbf{f}_n\}_{n=1, \dots, N-1}$ is the canonical base of \mathbb{R}^{N-1} . We define:

$$\mathbf{h}_n = \frac{\sum_{i=1}^n \mathbf{e}_i - (n-1)\mathbf{e}_{n+1}}{\|\sum_{i=1}^n \mathbf{e}_i - (n-1)\mathbf{e}_{n+1}\|},$$

with $\{\mathbf{e}_n\}_{n=1, \dots, N}$ the canonical base of \mathbb{R}^N . Note that $\phi(\mathcal{S}^{N-2}) = \mathcal{S}^{N-2}$, and that is why we can work in \mathbb{R}^{N-1} and then return in \mathbb{R}^N .

2.2 The AcSel algorithm

To use the selection-boost algorithm, we need a grouping method gr_{c_0} depending on an user-provided constant $0 \leq c_0 \leq 1$. This constant determines the strength of the grouping effect. The grouping method maps each variable index $1, \dots, P$ to an element of $\mathcal{P}(\{1, \dots, P\})$ (with $\mathcal{P}(S)$ is the powerset of the set S , *i.e.* the set which contains all the subsets of S). In practical terms, $gr_{c_0}(p)$ is the ensemble of all variables which are considered to be correlated to the variable \mathbf{x}_p and $\mathbf{X}_{gr_{c_0}(p)}$ is the submatrix of \mathbf{X} containing the row which indices are in $gr_{c_0}(p)$. If gr_{c_0} is looked as function which depends on c_0 it should have the following properties :

- $c_0 \in [0, 1]$,
- $gr_1(p)$ is the the set of all the indices of variables perfectly correlated (positively or negatively) to variable p ,
- $gr_0(p)$ is the set containing the indices of all the variables,
- if $c < c'$ then $gr_{c'} \subset gr_c$.

Furthermore, we need to have a selection method $select : \mathbb{R}^{N \times P} \times \mathbb{R}^N \rightarrow \{0, 1\}^P$ which maps the design matrix \mathbf{X} and the response variable \mathbf{y} to a 0-1 vector of length P with 1 at position p if the method selects the variable p and 0 otherwise.

Here, we make the assumption that a group of correlated variables are independent realizations of the same multivariate Gaussian law. As the variables are normalized with respect to the \mathcal{L}_2 norm, we will use the von-Mises Fisher law [Sra12] in \mathbb{R}^{N-1} thanks to the isomorphism ϕ . The probability density function of the von Mises-Fisher distribution for the random P -dimensional unit vector \mathbf{x} is given by:

$$f_P(\mathbf{x}; \boldsymbol{\mu}, \kappa) = \tilde{K}_P(\kappa) \exp(\kappa \boldsymbol{\mu}' \mathbf{x})$$

where $\kappa \geq 0$, $\boldsymbol{\mu} = (\mu_1, \dots, \mu_P)'$, $\|\boldsymbol{\mu}\|_2 = 1$ and the normalization constant $C_P(\kappa)$ is equal to

$$\tilde{K}_P(\kappa) = \frac{\kappa^{P/2-1}}{(2\pi)^{P/2} I_{P/2-1}(\kappa)},$$

where I_v denotes the modified Bessel function of the first kind and order v [AS72].

We then use the von-Mises Fisher law to create replacement of the original variables by some simulations (see Algorithm 1) to create B new design matrices $\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(B)}$. The AcSel algorithm then applies the variable selection method *select* to each of these matrices and returns a vector of length P with the frequency of apparition of each variable. The frequency of apparition of variable \mathbf{x}_p , noted ζ_p is assumed to be an estimator of the probability $\mathbb{P}(\mathbf{x}_p \in \mathcal{S})$ for this variable to be in \mathcal{S} . Nevertheless, both the grouping method and the choice of c_0 are crucial. When this constant is too small, the model is not enough perturbed. On the other hand, when this constant is too large, variables are chosen at random.

Algorithm 1 Pseudo-code for the AcSel algorithm with c_0 fixed

Require: $gr_{c_0}, select, B, c_0, P$

```

 $\zeta \leftarrow \mathbf{0}_P$ 
for  $b = 1, \dots, B$  do
   $\mathbf{X}^{(b)} \leftarrow \mathbf{X}$ 
  for  $p = 1, \dots, P$  do
     $\mathbf{x}_p^{(b)} \leftarrow \phi^{-1} \left( \text{random-vMF} \left( \hat{\boldsymbol{\mu}}(\phi(\mathbf{X}_{gr_{c_0}(p)})), \hat{\kappa}(\phi(\mathbf{X}_{gr_{c_0}(p)})) \right) \right)$ 
  end for
   $\zeta \leftarrow \zeta + select(\mathbf{X}^{(b)}, \mathbf{y})$ 
end for
 $\zeta \leftarrow \zeta / B$ 

```

The AcSel algorithm returns the vector $\zeta = (\zeta_1, \dots, \zeta_P)'$. One has now to choose a threshold to determine which variables are selected. In this article, we choose to select a variable p if $\zeta_p = 1$. In some applications, lower choices of threshold may be chosen.

2.3 Choosing c_0

As we will show in the next session, the smaller the c_0 parameter, the higher the precision of the resulting selected variables. On the other hand, it is obvious that the probability of choosing none of the variables (*i.e.* resulting in the choice of the empty set) increases as the parameter c_0 decreases. In the perspective of experimental planning, the choice of c_0 should result of a compromise between precision and proportion of empty models. Nevertheless, the c_0 parameter can be used to introduce a confidence indicator γ_p related to the variable \mathbf{x}_p :

$$\gamma_p = 1 - \min_{\mathbf{x}_p \in \hat{\mathcal{S}}_{c_0}} c_0. \quad (8)$$

When γ_p is near to zero, we know that a little perturbation is enough for variable p not being in the set of selected variables. On the other hand, when γ_p is near to one, we know that a really strong perturbation is needed for variable p not being in the set of selected variables. Notice that

using the confidence indicator implies using the AcSel method with all $c_0 \in [0, 1]$, or at least with a discretized grid of $[0, 1]$.

2.4 Choosing the grouping method

The simplest method for the grouping function gr_{c_0} is the following:

$$gr_{c_0}(p) = \left\{ p' \in \{1, \dots, P\} \mid | < \mathbf{x}_p, \mathbf{x}_{p'} > | \geq c_0 \right\}. \quad (9)$$

In other words, the correlation group of the variable p is determined by variables whose correlation with \mathbf{x}_p is at least c_0 . In the following this method will be referred as the "naïve" grouping method. Nevertheless, the structure of correlation may further be taken into account using graph community clustering. Let \mathbf{C} be the correlation matrix of matrix \mathbf{X} . Let define $\check{\mathbf{C}}$ as follows:

$$\check{c}_{ij} = \begin{cases} |\check{c}_{ij}| & \text{if } |\check{c}_{ij}| > c_0 \text{ and } i \neq j \\ 0 & \text{otherwise} \end{cases}.$$

Then, we apply a community clustering algorithm on the undirected network with weighted adjacency matrix defined by $\check{\mathbf{C}}$.

2.5 Conclusion

In this section we introduced the AcSel algorithm. As reported, this method is dependant on three elements which are the initial selection function *select*, the grouping method gr_{c_0} , and the constant c_0 . Whereas any of the variable selection functions reviewed in the introduction can be used for the selection function, we provide two specific grouping methods. Furthermore, the next section will show that the choice of c_0 should be made in respect to practical considerations.

3 Numerical studies

3.1 Introduction

To access the performances of the AcSel algorithm, we performed numerical studies. As stated before, the AcSel algorithm can be applied to any existing variable selection method. Here, we decided to use the Lasso and forward stepwise selection. The performance of the Lasso is known to be strongly dependant on the choice of the penalty parameter λ . In our simulations, we used four criteria to choose this penalty parameter: BIC, modified BIC (BIC2) in which the estimation of the residual variance is calculated with the model including two variables, AICc which is known to be asymptotically equivalent to cross-validation and generalized cross-validation (GCV).

To demonstrate the performance of the AcSel method, we compared our method with stability selection and with a naïve version of our algorithm, naïveAcSel. The naïveAcSel algorithm works as follows: estimate β with any variable selection method then if $gr_{c_0}(p)$, as defined in equation (9) for example, is not reduced to $\{p\}$, shrink $\hat{\beta}_p$ to 0. The naïveAcSel algorithm is similar to the AcSel algorithm, except that it does not take into account the error which is made choosing at random a variable among a set of correlated variables.

We explored four situations. Let P be the number of variables and N the number of observations. Data are generated from model in equation (1), assuming that $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$. The variance σ^2 is chosen to reach a signal to noise ratio of 5. Exception made of situation 4, variables are simulated following a multivariate Gaussian law, with variance-covariance matrix Σ . The diagonal elements of Σ are always set to 1. Each situation is repeated 200 times.

Situation 1 We are in the case where $P = N = 10$ and $\beta = (1, 1, 1, 0, 0, \dots, 0)'$. We set $\Sigma_{ij} = 0$ for $1 \leq i \neq j \leq 9$ and $\Sigma_{1,10} = 0$.

Situation 2 We are in the case where $P = 50$ and $N = 20$ and $\beta = (1, 1, 1, 1, 1, 0, 0, \dots, 0)'$. We set $\Sigma_{ij} = 0.5$ for $1 \leq i \neq j \leq 50$.

Situation 3 We are in the case where $P = 500$ and $N = 25$ and $\beta = (1, 1, 1, 1, 1, 0, 0, \dots, 0)'$. We set $\Sigma_{ij} = 0.5$ for $1 \leq i \neq j \leq 500$.

Situation 4 In this situation we use gene expression from a microarray data experiment in which $N = 24$. We first select the 1300 genes that were differentially expressed (stimulated versus unstimulated). For each repetition, we randomly select 100 genes among the 1300 and use the model in equation (1) to generate the response variable. We set $\beta = (1, 1, 1, 1, 1, 0, 0, \dots, 0)'$.

We use 4 indicators to evaluate the abilities of our method on simulated data. We define:

- recall as the ratio of the number of correctly identified variables (*i.e.* $\hat{\beta}_i \neq 0$ and $\beta_i \neq 0$) over the number of variables that should have been discovered (*i.e.* $\beta_i \neq 0$).
- precision as the ratio of correctly identified variables (*i.e.* $\hat{\beta}_i \neq 0$ and $\beta_i \neq 0$) over the number of identified variables (*i.e.* $\hat{\beta}_i \neq 0$).
- Fscore as the following ratio:

$$2 \times \frac{\text{recall} \times \text{precision}}{\text{recall} + \text{precision}}.$$

- emptiness as the proportion of empty models (no variable is selected)

Recall, precision, and Fscore are calculated over all models that are not empty. Note that our interest is focused on precision, as our goal is to select reliable variables. When $c_0 = 1$ the AcSel algorithm has no difference with the initially selected method *select*. When c_0 is decreasing toward zero we expect a profit in precision and a decrease of recall. We also calculate the Fscore which combines both recall and precision. As an improvement of precision comes with an increase of the proportion of empty models, the best method is one with the highest precision for a given level of emptiness.

3.2 Results of the simulation

Only an extract of the results is presented in the main part of the article; full results are available in supporting informations. We first analyse the results for each pair of selection method and situation. We show the evolution of the four criteria (precision, recall, Fscore and emptiness) in

function of the decrease of c_0 . When $c_0 = 1$, the AcSel algorithm is equivalent to the initial variable selection method. As our main focus is on precision, we add three histograms representing the evolution of the precision distribution for the highest, an intermediate and the lowest c_0 . Figure 2 shows the result for the Lasso with the modified BIC in Situation 1 (other Situations for BIC2 can be found in supplementary figures C.1, C.2, C.3 and C.4). In this example, we succeed to improve precision from 0.63 to 0.93. Other variable selection methods show interesting improvement of precision: the gain in precision is the lowest for the Lasso with the BIC criterion. This is not surprising since this method reaches the highest level of precision when $c_0 = 1$. On the other hand, the Lasso with AICc or GCV (see Figures A.1, A.2, A.3, A.4) present the greatest improvement in precision with the decrease of c_0 : in Situation 2, for the Lasso with GCV, precision improves from 0.25 to 0.75. However, as shown by the histograms of the precision, the proportion of models for which precision reaches one increases with the decrease of c_0 . The Fscore remains either stable or shows a small decrease indicating that the loss in recall is compensated by the increase of precision. In other words, our method allows to choose the desired trade-off between recall and precision (see Figures A.1, A.2, A.3, A.4).

As our interest is focused on precision, our goal is to reach the highest precision with the lowest proportion of empty models. In this context, one interesting fact about the AcSel method is that the method of choice of the penalty parameter in the Lasso is no more crucial. Indeed, as shown in Figure 3, the precision of each method is similar at a given proportion of emptiness. Nevertheless, depending of the situation, the choice of the penalty parameter by AICc (see Figures A.1, A.2, A.3, A.4) or GCV (see Figure 4) may lead to worse outcomes, even if there is an increase of precision with the confidence index.

Except in one case (the Lasso with choice of penalty parameter through the BIC criterion, see Figure 3), the AcSel algorithm shows its superiority over the naiveAcSel algorithm. The error which is made when choosing randomly a variable among a set of correlated variables conduces to further wrong choice of variables. While the intensive simulation of our algorithm allows to take into account this error, the naiveAcSel does not. The superiority of the naive algorithm in Situation 1 for the Lasso with BIC criterion may be the consequence of the small size of the data and the low correlation setting. Other situations are shown in Figures ??, ??, ?? and ??.

Finally we compare the AcSel algorithm with stability selection. Stability selection use a re-sampling algorithm to determine which of the variable which are included in the model are robust. In our simulation, stability selection shows performances with relative high precision but also high proportion of empty models. Moreover, in contrast to the AcSel algorithm, stability selection does not allow to choose a convenient precision-emptiness trade-off.

In the previous section, we mentioned the possibility of using AcSel to obtain a confidence indicator, corresponding to one minus the lowest c_0 for which a variable is selected. For each situation, we plot the proportion of correctly identified variables in function of the confidence indicator (Figure 4 for Situation 1 and Supplemental Figures for the others). The proportion of correctly identified variables increases with the increase of the confidence indicator.

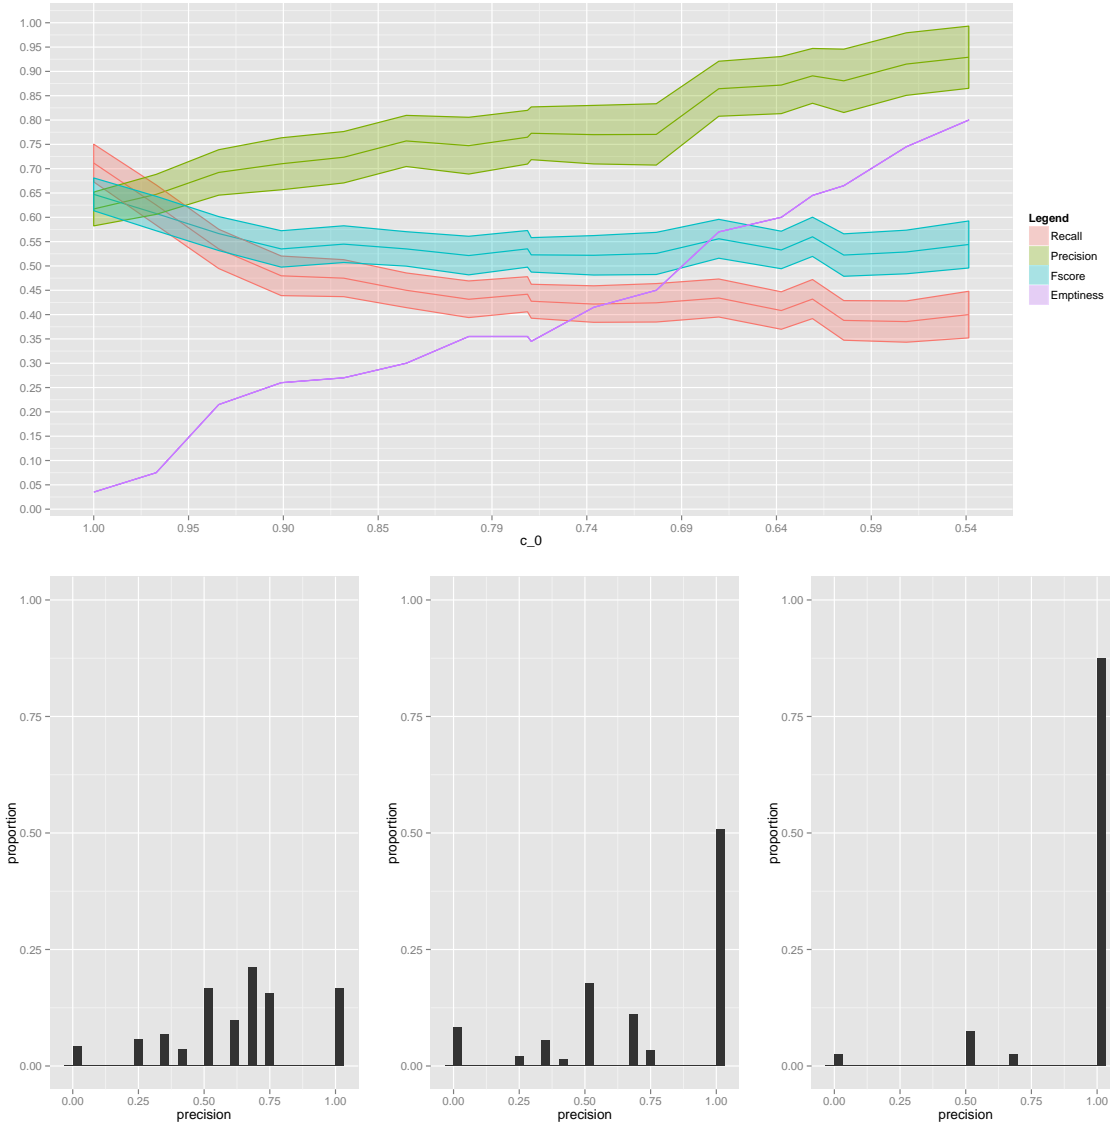


Figure 2: Example of result, here Situation 1 with the Lasso with the modified BIC criterion. Top figure: evolution of the four indicators (recall, precision, Fscore and emptiness) with 95% bandwidth confidence interval in function of c_0 . Bottom: the distribution of the precision among all non-empty models for the highest, an intermediate, and the lowest c_0 .

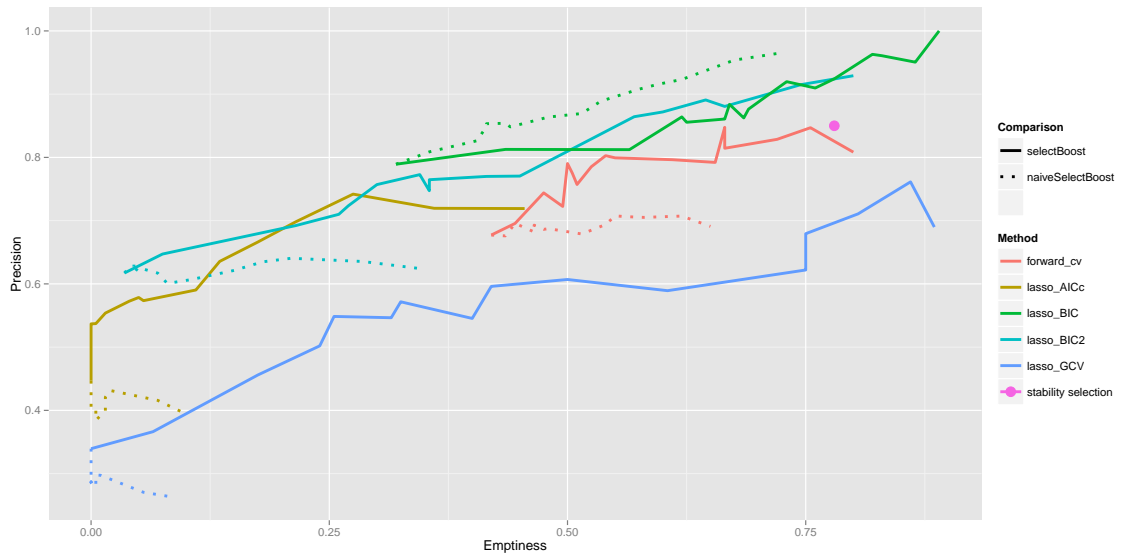


Figure 3: Precision in function of emptiness for all tested method for Situation 1. The AcSel algorithm is compared to both stability selection and the naiveAcSel algorithm.

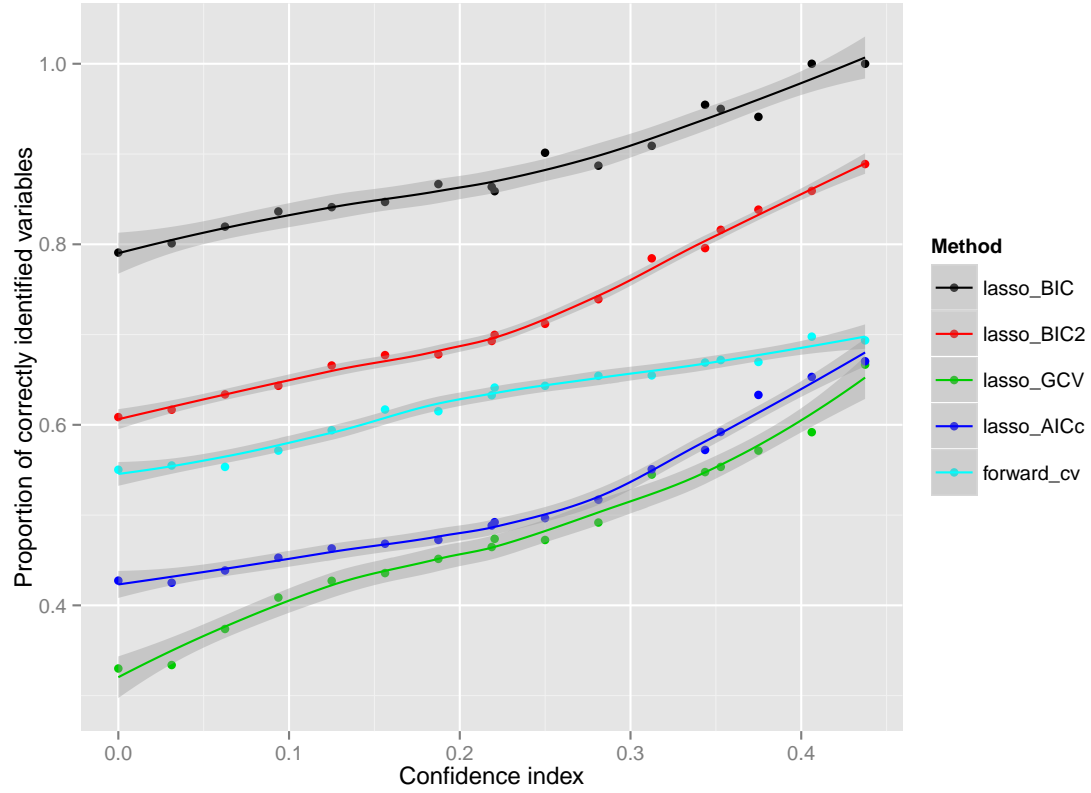


Figure 4: The proportion of correctly identified variables is plotted in function of the confidence indicator that we defined in Equation (6.8) for Situation 1. As expected, the greater the confidence indicator, the higher the proportion of correctly identified variables.

4 Application to a real dataset

We decided to apply our algorithm to the diabetes dataset used by Efron *et al.* [EHJT04]. This dataset contains 10 variables which are age, sex, body mass index, average blood pressure and six serum measurements and a quantitative response of interest that is a measure of the evolution of the diabetes disease one year after baseline. As proposed, we included interactions term, resulting in an 64 explanatory variables dataset.

We use a wide range of the c_0 parameter, starting from 1 to 0.35 by step of 0.05 (see Figure 5 right). For each step, the probability of being included in the support is calculated with 500 simulations as described in the Algorithm 1. We set the threshold to 0.95 to avoid numerical instability. We used our algorithm with the Lasso and selected the regularization parameter with the AICc.

When $c_0 = 1$, our algorithm is equivalent to the Lasso and ends with a selection of 22 variables. At the opposite, when using the maximal x_0 , our algorithm ends with a selection of only two variables : the body mass index and the average blood pressure. The interesting point is that these two variables are neither the two first variables selected by the Lasso or the two variables with the highest coefficients (see Figure 5 left). This demonstrates that our algorithm can be very useful to determine which are the variables that are selected with confidence.

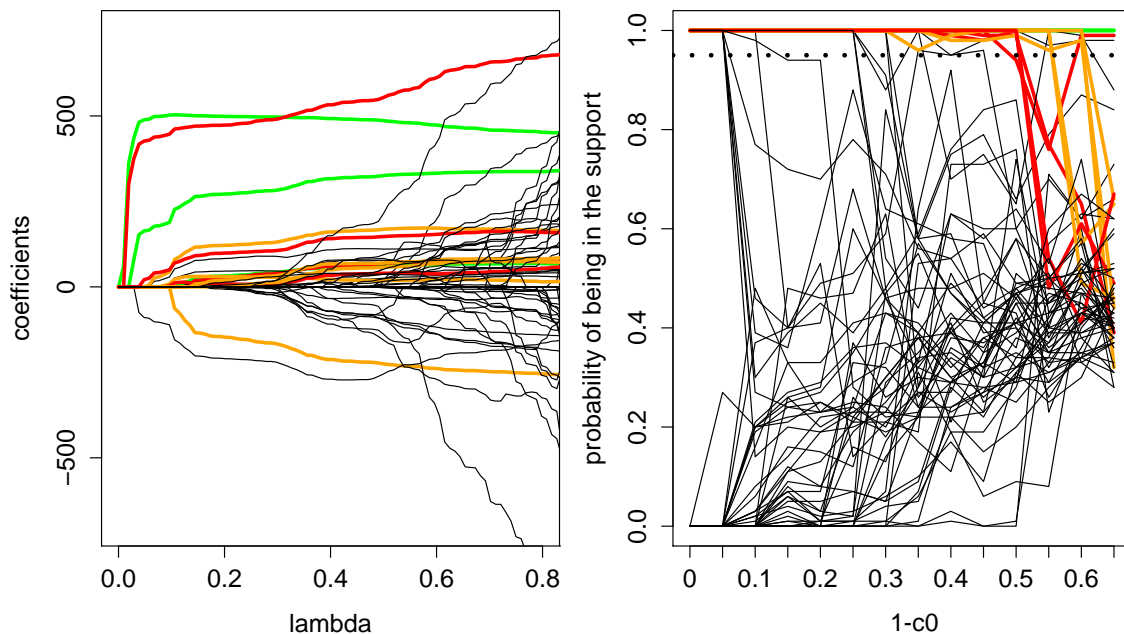


Figure 5: Colors: the green is for the most reliable variables selected by the AcSel algorithm (confidence indicator of 0.65; orange is for intermediate confidence (0.55) and red for low confidence (0.45). Right: evolution of the coefficients in the lasso regression when the sparsity parameter λ is varying. right: evolution of the probability of being in the support of the regression when the confidence indicator is varying. For the two graphics, each line maps with a variable.

5 Discussion and conclusion

We introduced the AcSel algorithm which uses intensive computation to select variables with high precision. The user has the choice between using this algorithm to produce a confidence indicator, or choosing an appropriate precision-emptiness trade-off to select variables with confidence. The main idea of our algorithm is to take into account the correlation structure of the data and thus use intensive computation to select the reliable variables. We prove the performance of our algorithm through simulation studies. In some situations, where many regressions have to be made (in network reverse-engineering in which we have a regression per vertex) our algorithm may be used in an experimental design approach. We apply our algorithm on a real dataset and we found some interesting informations. The AcSel algorithm is a powerful tool that can be used in every situation where reliable and robust variable selection has to be made.

References

- [Aka74] Hirotugu Akaike. A new look at the statistical model identification. *Automatic Control, IEEE Transactions on*, 19(6):716–723, 1974.
- [AS72] Milton Abramowitz and Irene A Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables. National Bureau of Standards Applied Mathematics Series 55. Tenth Printing*. ERIC, 1972.
- [BA02] Kenneth P Burnham and David R Anderson. *Model selection and multi-model inference: a practical information-theoretic approach*. Springer, 2002.
- [Bac08] Francis R Bach. Bolasso: model consistent lasso estimation through the bootstrap. In *Proceedings of the 25th international conference on Machine learning*, pages 33–40. ACM, 2008.
- [CDS01] Scott Shaoping Chen, David L Donoho, and Michael A Saunders. Atomic decomposition by basis pursuit. *SIAM review*, 43(1):129–159, 2001.
- [CGHE07] Li Chen, Dmitry B Goldgof, Lawrence O Hall, and Steven A Eschrich. Noise-based feature perturbation as a selection method for microarray data. In *Bioinformatics Research and Applications*, pages 237–247. Springer, 2007.
- [CS94] John R Cook and Leonard A Stefanski. Simulation-extrapolation estimation in parametric measurement error models. *Journal of the American Statistical Association*, 89(428):1314–1328, 1994.
- [D⁺00] David L Donoho et al. High-dimensional data analysis: The curses and blessings of dimensionality. *AMS Math Challenges Lecture*, pages 1–32, 2000.
- [DE03] David L Donoho and Michael Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization. *Proceedings of the National Academy of Sciences*, 100(5):2197–2202, 2003.
- [EHJT04] Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.

- [EZ12] Martin Eklund and Silvelyn Zwanzig. Simsel: a new simulation method for variable selection. *Journal of Statistical Computation and Simulation*, 82(4):515–527, 2012.
- [Fan97] Jianqing Fan. Comments on «wavelets in statistics: A review» by a. antoniadis. *Journal of the Italian Statistical Society*, 6(2):131–138, 1997.
- [FF93] LLdiko E Frank and Jerome H Friedman. A statistical view of some chemometrics regression tools. *Technometrics*, 35(2):109–135, 1993.
- [FG94] Dean P Foster and Edward I George. The risk inflation criterion for multiple regression. *The Annals of Statistics*, pages 1947–1975, 1994.
- [FHT10] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. A note on the group lasso and a sparse group lasso. *arXiv preprint arXiv:1001.0736*, 2010.
- [FL01] Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 2001.
- [FL06] Jianqing Fan and Runze Li. Statistical challenges with high dimensionality: Feature selection in knowledge discovery. *arXiv preprint math/0602133*, 2006.
- [FL10] Jianqing Fan and Jinchi Lv. A selective overview of variable selection in high dimensional feature space. *Statistica Sinica*, 20(1):101, 2010.
- [HK70] Arthur E Hoerl and Robert W Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- [Hoc76] Ronald R Hocking. A biometrics invited paper. the analysis and selection of variables in linear regression. *Biometrics*, 32(1):1–49, 1976.
- [KBIS99] John R Koza, Forrest H Bennett III, and Oscar Stiffelman. *Genetic programming as a Darwinian invention machine*. Springer, 1999.
- [LFGL99] Robert J Lipshutz, Stephen PA Fodor, Thomas R Gingeras, and David J Lockhart. High density synthetic oligonucleotide arrays. *Nature genetics*, 21:20–24, 1999.
- [LSB06] Xiaohui Luo, Leonard A Stefanski, and Dennis D Boos. Tuning variable selection procedures by adding noise. *Technometrics*, 48(2):165–175, 2006.
- [Mal73] Colin L Mallows. Some comments on cp. *Technometrics*, 15(4):661–675, 1973.
- [MB10] Nicolai Meinshausen and Peter Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):417–473, 2010.
- [Nat95] Balas Kausik Natarajan. Sparse approximate solutions to linear systems. *SIAM journal on computing*, 24(2):227–234, 1995.
- [Sch78] Gideon Schwarz. Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464, 1978.
- [SDC03] Mark R Segal, Kam D Dahlquist, and Bruce R Conklin. Regression approaches for microarray data analysis. *Journal of Computational Biology*, 10(6):961–980, 2003.

- [Sra12] Suvrit Sra. A short note on parameter approximation for von mises-fisher distributions: and a fast implementation of $\text{is}(\mathbf{x})$. *Computational Statistics*, 27(1):177–190, 2012.
- [Tib96] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- [WBS07] Yujun Wu, Dennis D Boos, and Leonard A Stefanski. Controlling variable selection by the addition of pseudovariables. *Journal of the American Statistical Association*, 102(477), 2007.
- [WL08] Tong Tong Wu and Kenneth Lange. Coordinate descent algorithms for lasso penalized regression. *The Annals of Applied Statistics*, pages 224–244, 2008.
- [WNRZ11] Sijian Wang, Bin Nan, Saharon Rosset, and Ji Zhu. Random lasso. *The annals of applied statistics*, 5(1):468, 2011.
- [YL06] Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- [ZH05] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.
- [Zha10] Cun-Hui Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894–942, 2010.
- [Zou06a] Hui Zou. The adaptive lasso and its oracle properties. *Journal of the American statistical association*, 101(476):1418–1429, 2006.
- [Zou06b] Hui Zou. The adaptive lasso and its oracle properties. *Journal of the American statistical association*, 101(476):1418–1429, 2006.
- [ZY06] Peng Zhao and Bin Yu. On model selection consistency of lasso. *The Journal of Machine Learning Research*, 7:2541–2563, 2006.

Supplementary information

A Confidence index

For the four simulation situations, we plot the proportion of correctly identified variables against the confidence index.

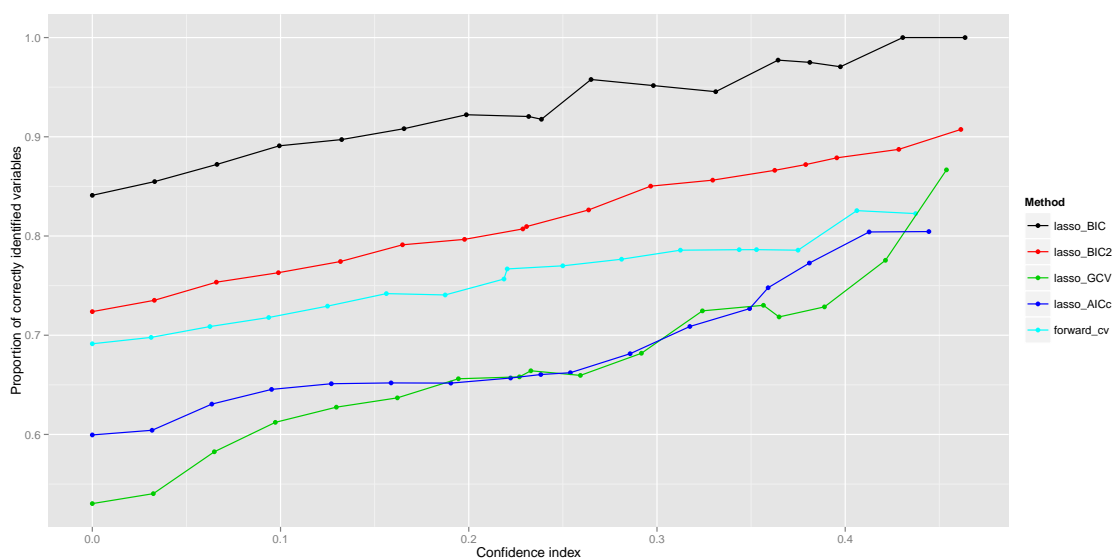


Figure A.1: Situation 1

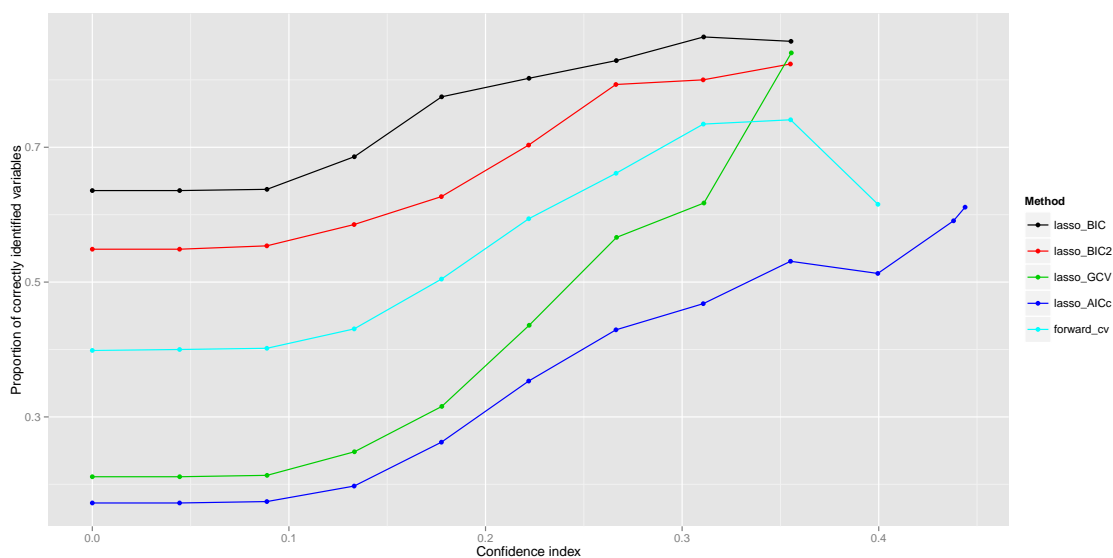


Figure A.2: Situation 2

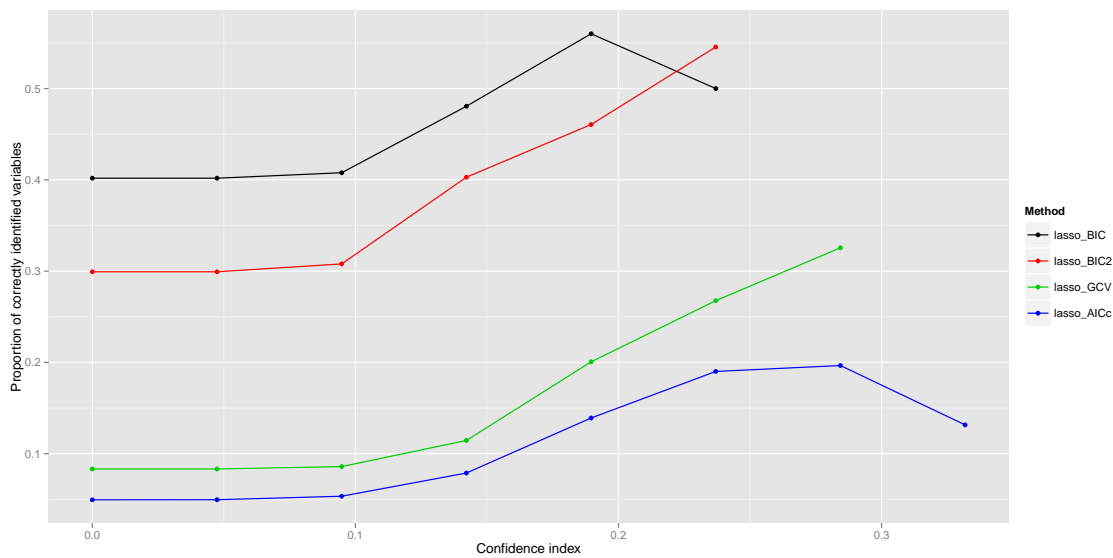


Figure A.3: Situation 3

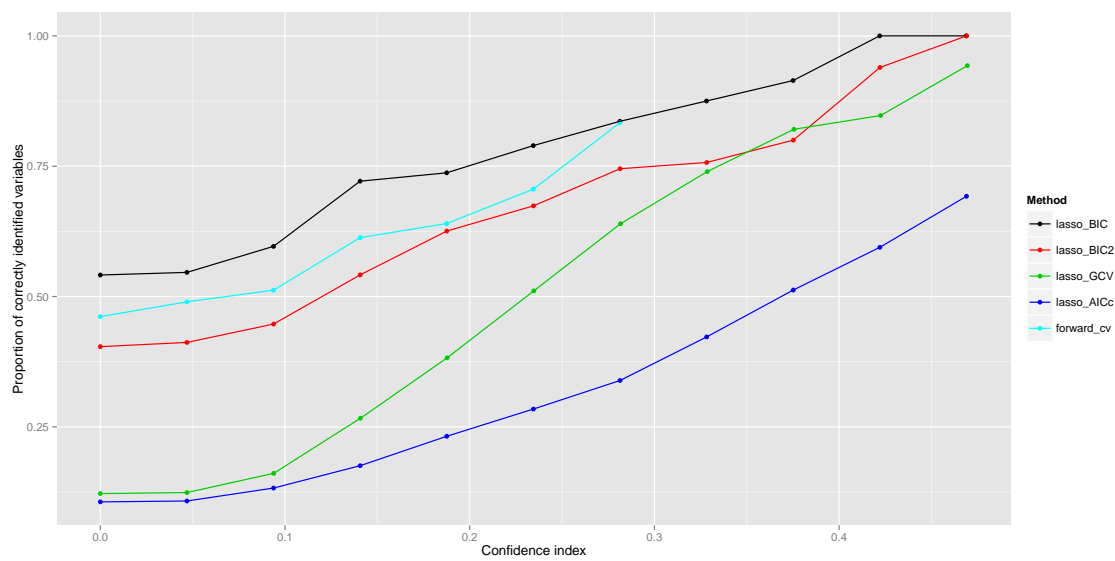


Figure A.4: Situation 4

B Comparisons

For the four simulation situations, we show the performance of our algorithm against the performance of the naive AcSel algorithm and the Stability Selection algorithm. The best algorithm is the one with the lowest proportion of empty models with the highest precision.

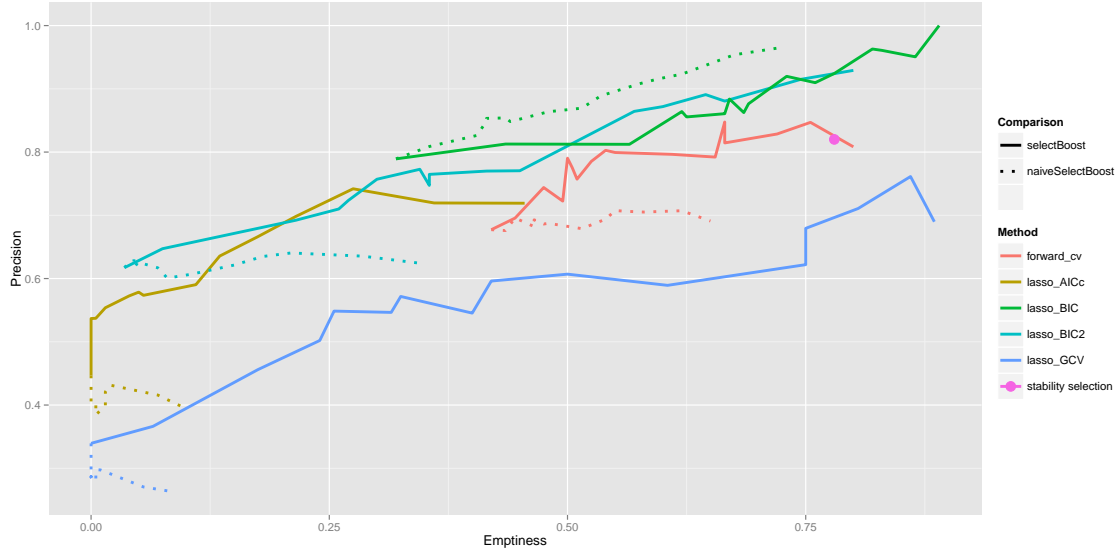


Figure B.1: Situation 1

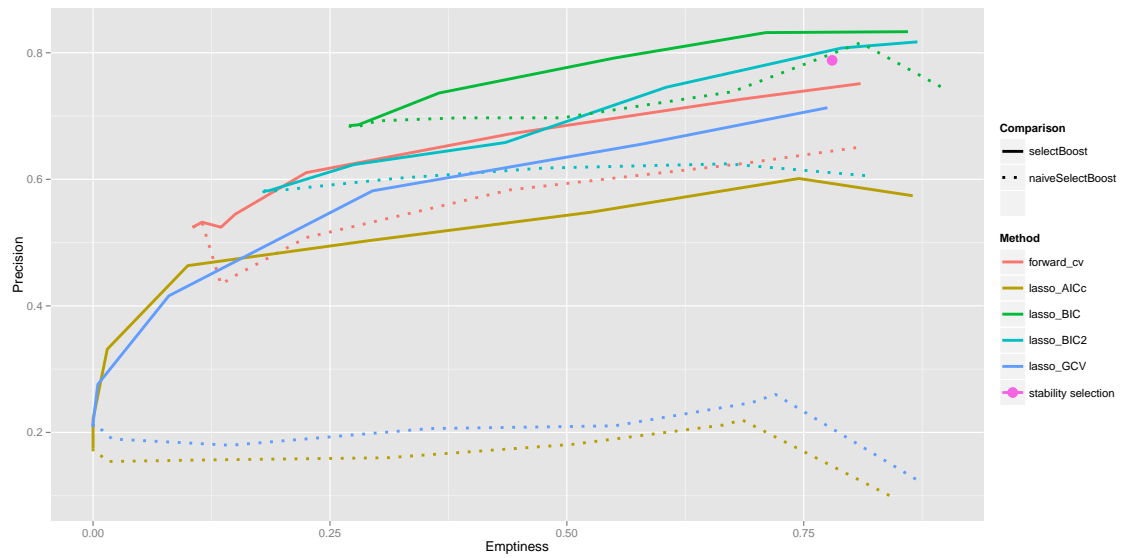


Figure B.2: Situation 2

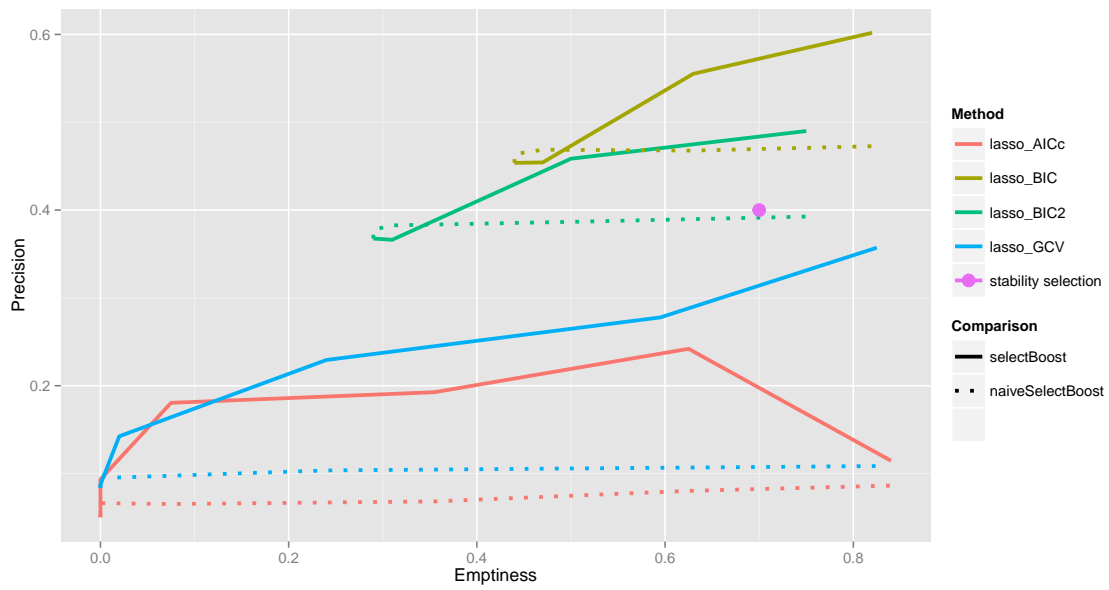


Figure B.3: Situation 3

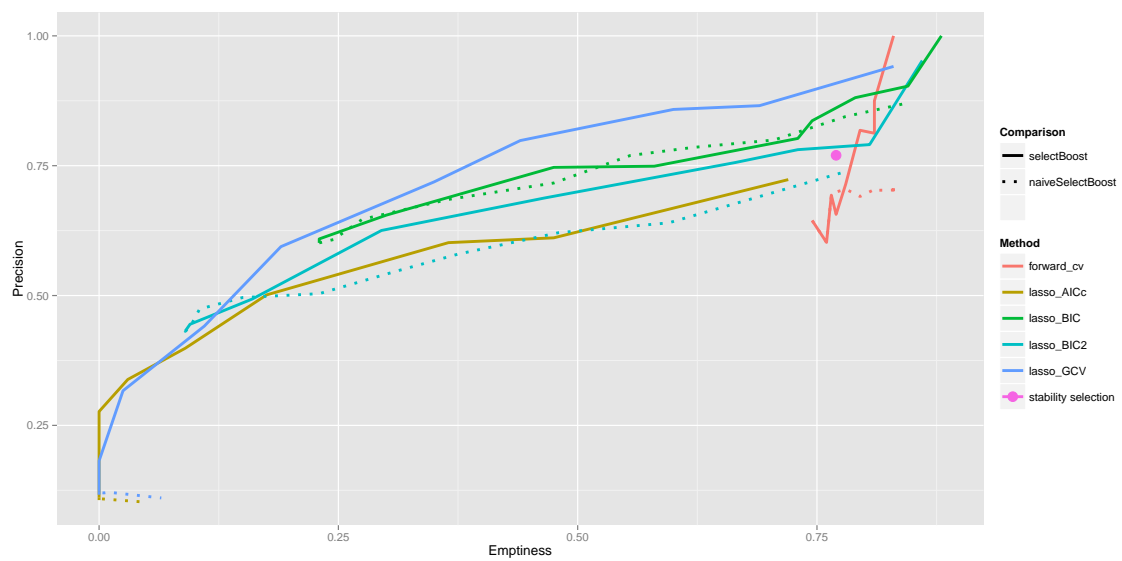


Figure B.4: Situation 4

C Example of results: modified BIC (BIC2)

In this section we show the evolution of recall, precision, Fscore and proportion of empty models with 95% confidence interval in function of the c_0 parameter. We also show three histograms with the evolution of the distribution of the precision for three c_0 (see legends)

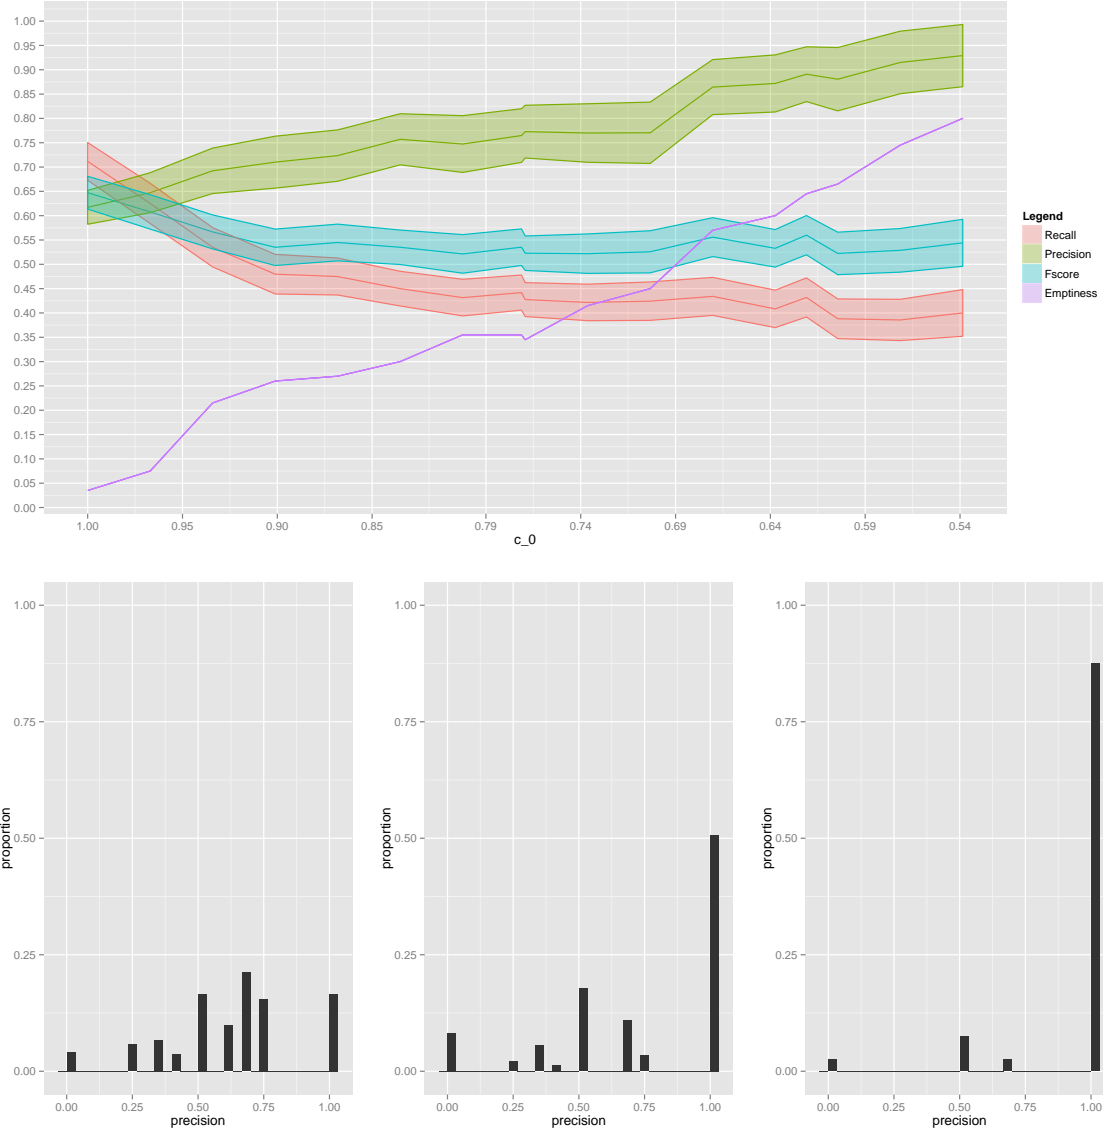


Figure C.1: Situation 1. The histograms show the evolution of the distribution of the precision. From left to right: $c_0 = 1, 0.79, 0.54$

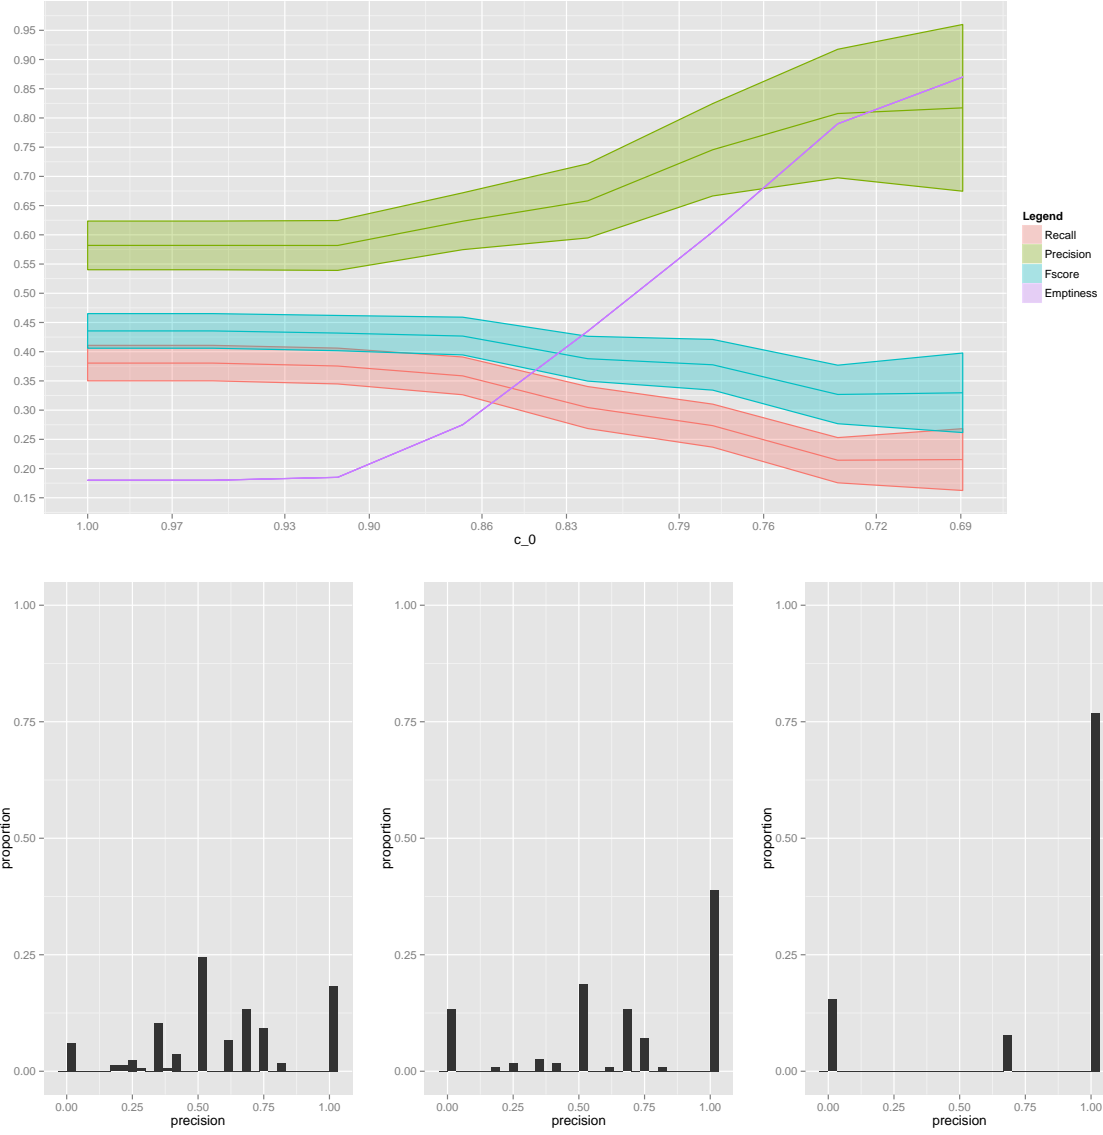


Figure C.2: Situation 2. The histograms show the evolution of the distribution of the precision. From left to right: $c_0 = 1, 0.79, 0.69$

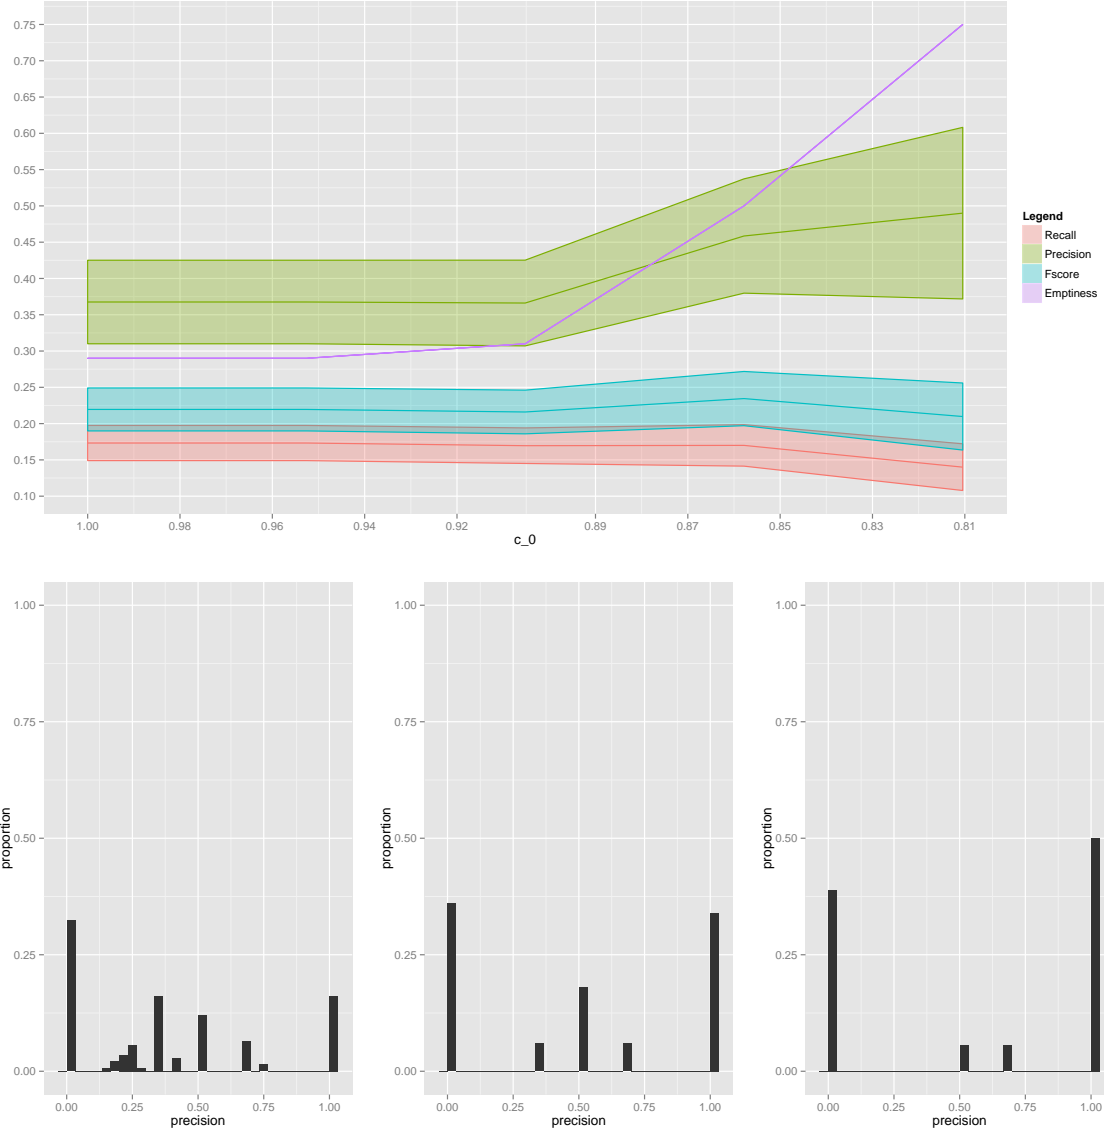


Figure C.3: Situation 3. The histograms show the evolution of the distribution of the precision. From left to right: $c_0 = 1, 0.9, 0.81$

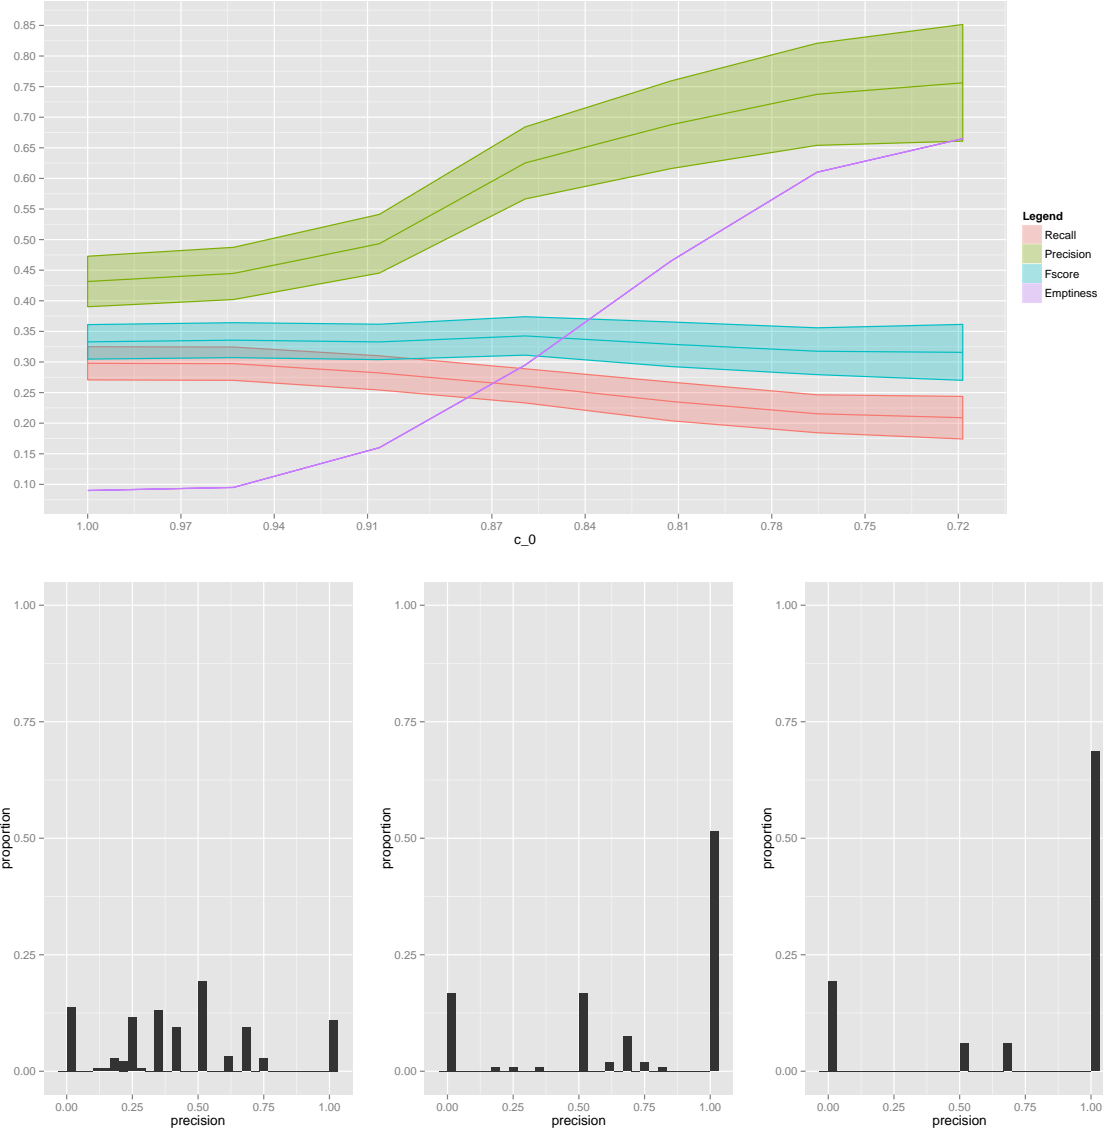


Figure C.4: Situation 4. The histograms show the evolution of the distribution of the precision. From left to right: $c_0 = 1, 0.87, 0.72$